LIBRARY
OF THE
MASSACHUSETTS INSTITUTE
OF TECHNOLOGY

HD28
.M414
no.92-64

WORKING PAPER
ALFRED P. SLOAN SCHOOL OF MANAGEMENT

A THEORY OF GROUP DECISION BEHAVIOR

92-64

Geoffrey P.E. Clarkson

and

Francis D. Tuggle

MASS. INST. TECH.
OCT 30 1964
DEWEY LIBRARY

MASSACHUSETTS
INSTITUTE OF TECHNOLOGY
50 MEMORIAL DRIVE
CAMBRIDGE, MASSACHUSETTS 02139

A THEORY OF GROUP DECISION BEHAVIOR

92-64

Geoffrey P.E. Clarkson

and

Francis D. Tuggle

MASS. INST. TECH.
OCT 30 1964
DEWEY LIBRARY

This paper is a draft for private circulation and comment. It should not be cited, quoted, or reproduced without the written consent of the authors. The research is supported by a grant from the Sloan School of Management from funds provided by NASA and the MAC-Project for research in organizational decision behavior.

HD28
.M414
no.92-

RECEIVED
NOV
M

A Theory of Group Decision Behavior

The investigation of group decision-making behavior raises a number of intriguing issues. Not the least of these are such questions as: Can a group's decision behavior with respect to a specific task be explained (predicted) from a knowledge of the decision processes of the individual participants? In what manner and by what mechanism does the process of arriving at a group decision affect the decision procedures of the individuals concerned? How and under what conditions do the psychological characteristics of each individual affect the group's decision-making behavior? Although there are many other questions which could be asked about the decision procedures of groups, this paper is principally concerned with presenting an answer to the first--to wit, a theory is proposed which from a knowledge of the decision processes of the individual members is sufficient to account for a group's decision behavior. However, while little attention is directed toward the remaining two questions, it is our opinion that the theory of group decision behavior as well as the method by which it is subjected to test provide a basis from which their answers can be readily developed. Accordingly, a discussion of these items is left until the end with the major part of the paper being devoted to the development and testing of the proposed theory of the group decision-making process.

Antecedents of the Theory

In order to explain the decision behavior which takes place within a group when it is engaged in a specified task, we posit that it is first necessary to know the decision processes with respect to this task of each member of the group. To know an individual's decision processes implies the existence of a theory of individual decision-making behavior. For such a theory is required if one

is to be able to specify in detail the decision processes of a particular individual. In our research on group decision behavior we have employed an information processing theory of human decision-making (Newell, Shaw, and Simon, 1958a). The object of this theory is to explain the process of human problem solving by identifying the types of decision processes humans employ while solving a variety of problems. It is a basic assumption of the theory that decision processes can be isolated as well as identified, and that they can be represented by whole programs of processing which are nothing more than a series of straight forward mechanical operations.

Since theories of individual behavior have been developed to explain a number of aspects of human information processing, e.g. rote learning (Feigenbaum, 1963), hypothesis testing behavior in a binary choice situation (Feldman, Tonge, and Kanter, 1961), and the acquisition of sequential pattern concepts (Simon and Kotovsky, 1963), there is sufficient evidence to suggest that decision-making behavior can be successfully studied in a number of empirical contexts. For instance, the decision behavior of individuals engaged in solution of problems in logic (Newell, Shaw, and Simon, 1957), geometry (Gelernter, Hansen, and Loveland, 1960), chess (Newell, Shaw, and Simon, 1958b), and portfolio selection (Clarkson, 1962), to mention but a few examples, can and have been used as the basis from which to test the empirical validity of many of the hypothesized decision processes. This is not to say that all the hypotheses of a particular theory of human decision behavior could be tested in this fashion. Manifestly, some hypotheses will be peculiar to specific problem contexts. The basic assumption, however, is that a certain number of these information processing hypotheses can be tested in a variety of situations,

and that this number is sufficient to guarantee the empirical testability of the resulting theory.

Implicit in this last assertion is the further assumption that invariances exist in the structure of the decision processes of different individuals. Indeed, it is assumed that these invariances not only exist but that theory can also be isolated, identified, and empirically confirmed. For example, the theory of human problem solving (Newell, Shaw, and Simon, 1958a) postulates the existence in an individual of a memory, some primitive information processes, and a hierarchy of decision rules. An application of this theory to a problem such as how humans acquire concepts of sequential patterns (Simon and Kotovsky, 1963), turns these postulates into testable hypotheses by specifying the contents of memory, the requisite information processes, as well as the content and order of the decision processes, such that the theory is now sufficient to account for the observed behavior. If it were not possible to identify the structure of these processes, then one could not transform these postulates into the testable hypotheses of a specific theory of human decision behavior (Clarkson and Pounds, 1963). In effect, it is assumed that unless structural invariances exist, like the ordering and processing of symbols in memory, among the decision processes of different problem solvers, then it is not possible to construct a testable theory of individual decision behavior in this fashion.

If it is possible, on the basis of the research already conducted, to accept the statement that theories can be constructed which explain the decision-making behavior of individual subjects, then the question immediately arises as to whether this theory can be extended to encompass the decision behavior of

groups. To develop a testable theory of individual behavior a postulate is employed which asserts the existence of certain structural invariances in the decision processes of all problem solvers. But groups of all sizes are composed of individuals. Hence, the ability to infer from individual to group decision behavior would be provided by a postulate which asserts the existence of invariances between the structure of individual and group decision processes.[1/] The basis of this postulate resides in inductive and empirical grounds--it cannot be proved as a theorem. Indeed the only way in which it can be supported, other than by an appeal to parsimony and its consistency with the first postulate of invariance, is through a series of empirical tests.

In order to subject the second postulate of invariance to empirical test an experiment is required where it is possible to demonstrate that a group's decision behavior can be explained in the same way as that of an individual. To provide the requisite data the experiment must be designed in such a way that the decision processes of both individuals and groups can be readily elicited. At the same time, we are interested in testing the hypothesis that the decision behavior of the group can be predicted from a knowledge of the decision processes of its individual members. Accordingly, the experiment must also enable the experimenter to ascertain the decision procedures of each subject with respect to the task prior to their taking part in a group decision.

[1/] In the remainder of the paper this posit will be referred to as the second postulate of invariance.

Once such an experiment is designed the one can then immediately subject the proposed theory to a series of empirical tests. For the experiment yields two sets of observations: those which are produced by the individuals as they perform the task by themselves, and those generated by the same individuals when they are behaving as a group on the same experimental task. Since the decision behavior of the group is to be interpreted within the same theoretical structure as that of the individuals, the success with which this representation is carried out will provide empirical support for the second postulate of invariance. Similarly, once the decision processes of each participant are identified the level of success with which they are employed to predict the decision behavior of the group will constitute a direct test of the empirical validity of the proposed theory.

Experiment and Procedures

The experimental task presented to each subject is that of making a series of bids, i.e. announcing a series of prices, in two separate markets.[2/] Subjects are given a set of instructions about the task which, in brief, contain the following points: A trial consists of making one bid, in terms of dollars and cents, in each of two, independent markets. The bids are to be competitive in that subjects are informed they are bidding against two markets which reflect

---

[2/] This task is adapted from some previous experimental research on the processes which individuals employ in order to decide which of several alternatives to attend to next (Pounds, 1964). The inquiry into the "problem solving control" process yielded sufficient evidence on the decision behavior of individuals in this experimental task to recommend its use in the current situation.

the bids of other people. A subject "wins" on his bids if his prices are below those generated by the market. Thus, at the end of each trial a subject is told whether he "won" or "lost" on his bids. These results are determined by the experimenter who compares the subject's bids to a particular series of random numbers.[3/] The task itself consists of a sequence of such trials where the object for a subject is to win as much money as he can. The only restriction which is imposed upon the bidding process is that a subject may only alter one of his bids on each trial. Accordingly, except for the first trial on which two bids are required,[4/] the problem faced by a subject is to decide which if either of his two previous bids to alter and in what direction (increase or decrease) so that he wins as much as he can over an announced series of such trials.

Before commencing to play a subject is presented with the written instructions and a special form upon which to write his bids. The form contains columns of numbered blanks, two per trial, into which the prices are entered. At the end of each trial the subject is instructed to draw a circle around those bids which won. Consequently, after a series of trials a subject has in front of him a complete record of his play, i.e. the number of times he has won or lost as well as the prices at which these events occurred. Manifestly, at the end

---

[3/] In these experiments sequences of random numbers are used which have a mean of 2 and a standard deviation of 1. In previous research (Pounds, 1964) various means and standard deviations were employed with no noticeable effect being observed in subjects' decision behavior.

[4/] The written instructions contain a suggested opening bid of $1 on each market. But as subjects are free to choose any price, many subjects ignore this suggestion and begin at higher prices.

of the announced set of trials this form constitutes a written record of a subject's bidding behavior in this task. And it is expressly this stream of behavior when generated by a group that we take as the events to be explained.

The experimental procedure is as follows: Subjects are provided with the instructions as well as a form and are told they are going to play 25 trials. These trials are used to permit subjects to gain experience with the task. In addition, they are requested during these trials to try and develop a rule or set of rules for behaving in this situation. After the 25th trial subjects are asked to write down the decision rules they were requested to think about as they played. To assist them in this process they are told to consider the problem in terms of trying to instruct someone else to play in their stead. They are also informed that as soon as they have finished writing their rules they will have 25 trials in which to test out its behavior. Accordingly, by the end of the 50th trial each subject will have had an opportunity to assess and amend, if he chooses, the set of rules by which he decides what prices to try next.

At this point, two subjects who have completed the first two stages are brought together in one room. They are informed that they are both going to play as a group for a third set of 25 trials. They still have their written rules as well as any comments they may have appended and they are given 10 to 15 minutes[5/] to discuss and play how they are going to play as a group. Once this decision

---

[5/] To date there has been almost no control on the amount of time allowed for the planning period. If after 15 minutes subjects have not agreed on a method for playing the experimenter suggests they might get on with it, but does not terminate the planning session. In some cases discussion was ended by the experimenter's suggestion that they try and work out what they are going to do as they go along--i.e. get started and see if conflicts cannot be resolved as they occur.

is made they are presented with a sequence of 25 trials the completion of which ends the experiment.

To recapitulate, the recorded decision behavior at the end of the experiment consists of the following items:

<u>Stage 1</u>: Bidding behavior in a set of 25 trials where a subject is learning how to play. These data are used by the subject as a basis for developing a rule or set of procedures for selecting bids.

<u>Stage 2</u>: A second set of 25 bids where the subject is trying out the rule derived from <u>Stage 1</u>.

<u>Stage 3</u>:[6/] A third set of 25 bids where two subjects are acting as a group to decide on what bids to select.

<u>The Theory and Definition of its Terms</u>

From the experimental procedure it is clear that we are interested in being able to explain the behavior observed during <u>Stage 3</u> from a knowledge of the individual behavior exhibited in <u>Stage 2</u>. Our theory is very simple and states: The behavior of a group is a direct consequent of the individual decision procedures with the addition of a process for resolving conflict (disagreement on bids) when it occurs. In other words, given the individual behavior expressed during <u>Stage 2</u> and a procedure for resolving conflict, a group's trial by trial decisions can be immediately explained.

[6/] The recorded behavior on all trials consists of two prices stated in terms of dollars and cents with the appropriate marks as to whether these bids won or lost. Also, it should be noted that the experimenter uses a different set of pairs of random numbers, drawn from the same population, for each of these stages.

To test this theory it is necessary that each term be operationally defined. To begin with, both an individual's and a group's decision behavior consists of that sequence of events which is recorded on their forms in Stages 2 and 3, with two slight modifications. These occur because we are primarily interested in explaining the sequence of bids in terms of the changes which take place from previous bids. In brief, we are not concerned with the actual prices themselves but rather with explaining the shifts in attention from market to market and the alterations in prices whether up or down. As a result, the actual behavior of an individual or a group is abstracted into a form which records trial by trial behavior in terms of wins, losses, increases, and decreases. If an individual or a group decides to bid one price on both markets for the entire 25 trials, there are no data to explain as no bidding changes have taken place even though there is a record of wins and losses. Accordingly, the decision behavior to be explained consists of a group's actual play described in terms of the attributes win, loss, increase, and decrease, with all sequences of two or more trials of constant price bids in both markets deleted.

For example, in Table 1 there is a series of five trials as recorded during Stage 3 with the corresponding abstracted sequence. It is the latter sequence, deleting the events marked "no change" that the theory is required to predict.

Stage 3		Abstracted Sequence
1.50	(1.75)	Loss - Win , IR
(1.50)	(1.85)	Win - Win , IL
(1.60)	1.85	Win - Loss , NC
1.60	1.85	Loss - Loss , DR
(1.60)	1.75	

I--Increase, D--Decrease, L--Left Market, R--Right Market, NC--No Change

Table 1

The second term to define is the expression "the individual decision procedures." There are presumably a number of ways in which these processes could be defined. One such is to take the rules and comments as written by the subjects before and during Stage 2 as the statement of what constitutes their decision rules. But, as they are at liberty to amend or completely alter their rules while playing this method has little to recommend it. Our approach is to take the data as produced by Stage 2 and process them through a Rule Generating Program. The function of this program is to infer the decision rules used by each subject to generate his bids. The Rule Generating Program's output is a decision rule which is stated in terms of the attributes listed in the first part of the second column of Table 1, e.g. Win - Win, Win - Loss, Loss - Win, Loss - Loss. Associated with each of these attributes are parameters which describe the behavior that results, in terms of increase, decrease, left, or right, when these events occur.

For instance, an individual who shifts from the left to the right market and back again every time he encounters a Win - Win situation would have an alternating side parameter associated with this attribute. The side parameter then keeps track of which market will be attended to next. Similarly, a subject who alternates between increases and decreases when faced with Loss - Win or Win - Loss situations would have an alternating direction parameter associated with one or both of these attributes. In some cases a subject's response to a particular situation, e.g. Loss - Win, depends on what happened on the preceding trial. In other cases, the subject behaves with admirable constancy such that his response to a given situation is always the same. Despite the variety of responses the Rule Generating Program is able to discover such

dependencies and include them in the description of a subject's decision rules. Accordingly, the output is a list of the four attributes with the relevant associated parameters. This list or set of rules is sufficient to account for the behavior recorded during Stage 2. In fact, the list constitutes precisely what we mean by the individual's decision procedures.

The third item to define is the process for resolving conflict. Conflict occurs when two individual decision rules produce different responses to the occurrence of a specific event, say a Loss - Loss situation. To resolve the conflict a third list of attributes or decision rules is used. These rules consist of a standard set of responses to the four possible events--Win- Win, Win - Loss, Loss - Win, Loss - Loss--with the addition of an alternating side parameter on the attributes Win - Win and Loss - Loss.[7/] These responses are given in Table 2, where S refers to the side (left or right market) specified by the side parameter.

Conflict Resolving Process

Situation	Response
Win - Win	IS
Win - Loss	DR
Loss- Win	DL
Loss- Loss	DS
S--Side indicated by the side parameter	

Table 2

7/ While there are a number of such processes which could function as conflict resolvers, the decision rule used here represents the actual process used by a number of groups. When pairs of subjects disagreed on what to do next they sometimes openly wrote down the rules given in Table 2 and used them as the basis upon which agreement on the next bid was reached.

To summarize, the theory states that a group's decision behavior is produced by the individual decision procedures plus a conflict resolving process. Consequently, to select a group's next bid the theory takes each of the two individual decision rules and notes their responses to the current situation. If they are the same, this response constitutes the theory's prediction. If they differ, then the theory's prediction is generated by the conflict resolving rule. Thus, the theory is able to produce a response for every situation and it is these responses which are the theory's predictions of a group's Stage 3 decision behavior.

Testing Procedures and Results

In order to subject the theory to empirical test five pairs of subjects were run through Stages 1-3 of the experiment.[8/] The results from Stage 2 processed by the Rule Generating Program and five pairs of individual decision rules were generated.[9/] Each pair of decision rules was then coded into the form required by the theory.[10/]

[8/] The only exception to the experimental procedure outlined above was that in some cases Stage 3 contained more than 25 trials. Our object was to collect as much data as possible and if time permitted more than 25 trials were run.

[9/] This program is currently in flow chart form only. Hence, the results presented here are derived by hand simulation. It is expected that the program will be coded and operating before long so that the entire theory can be executed by a computer.

[10/] This part of the theory is coded in Information Processing Language V, (Newell, et. al., 1961), and was run on an IBM 1620, Mark II.

To be able to predict the behavior of each group the theory must be provided with a similar history of experience. Hence, once the individual decision rules are derived the theory requires one additional piece of information--the first bids on both markets made by the group and whether they won or lost. Given this information the theory makes its prediction for the next trial. At the end of each trial it is then given the relevant history of Wins and Losses as experienced by this group in Stage 3. By the end of the group's trials the theory has generated a predicted sequence of behavior for all but the first. These predictions are directly compared with the abstracted sequence derived from the group's recorded decision behavior.[11/]

The theory's success in predicting the sequence of decisions made by these five groups is shown in Table 3. As noted above, the total number of trials during Stage 3 varied from group to group. All groups, however, played at least 25. Hence, the second column of Table 3 records the total number experienced by the group less those on which no change took place. The remaining columns contain the proportions of correct responses produced by the theory as well as statistical estimates of the likelihood of these proportions being achieved by chance--i.e. a process that selects responses at random.

In all cases the theory's predictions of the group's decision behavior is significantly different from a random process at the .01 level. On the

[11/] To facilitate this comparison the record of the group's behavior is also placed in the computer so that both the predicted and the observed responses can be printed out trial by trial.

selection of the correct side alone the results are less impressive. For there

Group	Total No. of Trials	Number Correctly Predicted	Total Proportion Correct	Level [12] of Significance	Proportion of Sides Chosen Correctly	Level [13] of Significance	Proportion of Direction Chosen Correctly	Level of Significance
1	14	10	.71	.01	.71		.93	.01
2	24	18	.75	.01	.75	.05	.96	.01
3	36	24	.67	.01	.67	.05	1.00	.01
4	37	33	.89	.01	.92	.01	.97	.01
5	16	9	.56	.01	.56		.94	.01

Table 3

are two cases, groups 1 and 5, where the theory's responses are not significantly different from what one would expect by chance. Further, of the remaining three groups, the results for two of them are significant at the .05 level while only the set of side predictions are significant at the .01 level. With respect to a group's choice of direction, however, the theory's predictions are uniformly significant at the .01 level.

---

[12] There are four possible attributes that can be selected at each trial, e.g. Increase Left, Increase Right, Decrease Left, Decrease Right. Therefore, by choosing at random one would expect to be correct 25% of the time. Hence, the level of significance of the observed proportion is computed by employing the normal approximation to the binomal distribution with $\mu=\hat{p}=.25$ and $\sigma=\sqrt{\hat{p}\,\hat{q}/n} = \sqrt{.25 \times .75/n}$.

[13] Both the side and the Direction can be expected to be chosen correctly by a random device 50% of the time. Thus, the level of significance of the observed proportion is computed by employing the normal approximation to the binomal distribution with $\mu=\hat{p}=.50$ and $\sigma=\sqrt{\hat{p}\,\hat{q}/n} = \sqrt{.50 \times .50/n}$. The blank spaces represent occasions where the result is not significant at the .05 level.

Implications for the Study of Group Behavior

The theory of group decision behavior as well as the results presented above demonstrate the possibility of being able to explain in some detail the decision-making process that takes place in groups. We freely admit that our theory has not yet been extensively tested and that the results reported here pertain solely to two person groups. However, we have begun to test a slightly amended version of the theory against the decision behavior of these person groups, and the results to date, while less striking than in the two person situation, are significant at the .01 level. Consequently, it appears as though an aspect of group decision behavior has been isolated as well as identified and that an information processing theory is sufficient to account for this decision-making behavior.

The theory itself, it should be noted, is endowed with no additional capabilities than are found in similar theories of human decision-making behavior (Simon and Kotovsky, 1963). For example, the theory postulates that each individual is capable of developing his decision rules in terms of the attributes of the experimental task, e.g. wins, losses, increases, decreases, prices, etc. The theory also posits that once an individual has constructed a rule he is capable of employing it whenever it is called for--i.e. activated by the present situation. In addition the theory is provided with the ability to notice both differences and similarities between symbols and is able to keep track of a small number of such symbols in immediate memory at one time. The remaining properties are described in an earlier section and need not be repeated. The important point is that a theory composed of these properties is sufficient to account for the decision behavior of groups as well as individuals. In brief,

the evidence suggests that the second postulate of invariance is empirically sound.

At the beginning of the paper three questions concerning group decision behavior are raised. The first concerns the possibility of explaining the decision behavior of a group from a knowledge of the decision processes of its members. Our theory and results appear to provide an affirmative answer to this question. Moreover, it is our contention that the theory also provides the empirical basis from which answers can be derived for the remaining two queries.

In order to discover how and by what mechanism the process of arriving at a group decision affects the individual's decision processes only one additional stage would be required in the experiment. Stage 4 would consist of requiring the individuals to play by themselves for another 25 trials. These data would then be processed by the Rule Generating Program which would yield a current description of their decision rules. The decision rules for each subject from Stages 2 and 4 could then be compared. Such a comparison assumed that any change in an individual's decision rule is a direct result of the group decision-making experience. To subject this assumption to test it would be necessary to be able to describe the process as well as the conditions under which such alterations take place. In short, a theory is required which accounts for the process by which individuals acquire and amend their decision procedures. While we are not at present able to produce such a theory, the inclusion of Stage 4 would appear to provide the requisite data from which it could be developed.

The third query concerns the effect on a group's decision behavior of the psychological characteristics of the participants. Although no theory yet exists

which links specific psychological characteristics to particular types of decision rules, it is not unreasonable to conjecture that such a relation might be empirically true. One method of investigating the merits of this hypothesis would be to present all subjects with a standard battery of tests prior to Stage 1. If a comparison of their psychological characteristics and their decision rules revealed significant correlations, then these results would tend to support the general proposition. However, in order to test a specific hypothesis concerning the relation between certain psychological characteristics and a particular class of decision rules it would also be necessary to be able to classify decision procedures into a number of different types. For if one is to be able to demonstrate that decision-making processes reflect the psychology of the subject, then it must be possible to recognize as well as categorize all differences in the former. While to our knowledge no research has yet been conducted on the categorization of individual decision processes, experiments are being conducted to determine the empirical merit of this conjecture.

References

Clarkson, G.P.E., Portfolio Selection: A Simulation of Trust Investment, Englewood Cliffs, New Jersey, Prentice-Hall, 1962.

Clarkson, G.P.E. and Pounds, W.F., "Theory and Method in the Exploration of Human Decision Behavior," Industrial Management Review, 5, 1963, pp. 17-27.

Feigenbaum, E.A., "The Simulation of Verbal Learning," in E.A. Feigenbaum and J. Feldman, (eds.), Computers and Thought, New York, McGraw-Hill, 1963, pp. 297-309.

Feldman, J., Tonge, F.M., and Kanter, H., "Empirical Explorations of a Hypothesis Testing Model of Binary Choice Behavior," Report No. SP-546, December 19, 1961, System Development Corporation, Santa Monica.

Gelernter, H.L., Hansen, J.R., and Loveland, D.W., "Empirical Explorations of the Geometry Theorem Machine," Proceedings of the Western Joint Computer Conference, 1960, pp. 145-159.

Newell, A., Shaw, J.C., and Simon, H.A., "Empirical Explorations of the Logic Theory Machine," Proceedings of the Western Joint Computer Conference, February, 1957, pp. 218-230.

________________, "Elements of a Theory of Human Problem Solving," Psychological Review, 65, 1958a, pp. 151-166.

________________, "Chess-Playing Programs and the Problem of Complexity," IBM Journal of Research and Development, October, 1958b, pp. 320-325.

Newell, A., et. al., Information Processing Language V Manual, Englewood Cliffs, New Jersey, Prentice-Hall, 1961.

Pounds, W.F., "A Study of Problem Solving Control," unpublished Ph.D. thesis, Carnegie Institute of Technology, 1964.

Simon, H.A. and Kotovsky, K., "Human Acquisition of Concepts for Sequential Patterns," Psychological Review, 70, 1963, pp. 534-546.

BASEMENT

Date Due

OCT 16 '75		
SEP 11 '78	[redacted]	
APR 10 '78	DEC 15 1986	
SEP 15 '79	JUL 10 1989	
FEB 07 '80		
JUL 12 '8		
[redacted]		
[redacted]		
[redacted]		
[redacted]		

Lib-26-67

3 9080 003 900 260 85-64

MIT LIBRARIES
3 9080 003 869 218 86-64

MIT LIBRARIES
3 9080 003 900 203 87-64

1751831
Loan

MIT LIBRARIES
3 9080 003 900 211 88-64

MIT LIBRARIES
3 9080 003 900 476 89-64

MIT LIBRARIES
3 9080 003 869 457 90-64

MIT LIBRARIES
3 9080 003 869 507 91-64

MIT LIBRARIES
3 9080 003 900 500 92-64

MIT LIBRARIES
3 9080 003 900 484 93-64

MIT LIBRARIES
3 9080 003 869 465 94-64

MIT LIBRARIES
3 9080 003 900 468 97-64

MIT LIBRARIES
3 9080 003 869 481 98-64

MIT LIBRARIES
3 9080 003 869 531 99-64

95

96